

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

Claims 1-4 (Canceled).

5. (Previously presented) In a data processing network including distributed processing units, a method comprising:

obtaining a respective utilization value of each distributed processing unit;

applying a function to the respective utilization value of said each distributed processing unit to obtain a respective weight for said each distributed processing unit; and

using the respective weights for the distributed processing units for distributing work requests to the distributed processing units so that the respective weight for said each distributed processing unit specifies a respective frequency at which the work requests are distributed to said each distributed processing unit;

wherein the respective weight for said each distributed processing unit is programmed into a mapping table, and the function is applied to the respective utilization value of said each distributed processing unit to obtain the respective weight for said each distributed processing unit by indexing the mapping table with the respective utilization value of said each distributed processing unit to obtain the respective weight for said each distributed processing unit.

Claims 6-9 (Canceled).

10. (Previously presented) In a data processing network including distributed processing units, a method comprising:

obtaining a respective utilization value of each distributed processing unit;

applying a function to the respective utilization value of said each distributed processing unit to obtain a respective weight for said each distributed processing unit; and

using the respective weights for the distributed processing units for distributing work requests to the distributed processing units so that the respective weight for said each distributed processing unit specifies a respective frequency at which the work requests are distributed to said each distributed processing unit;

wherein the respective weights for the distributed processing units are used for distributing work requests to the distributed processing units by creating a distribution list containing entries indicating the distributed processing units, the respective weight for said each distributed processing unit specifying the number of the entries indicating said each distributed processing unit, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the work requests to the distributed processing units as indicated by the entries in the randomized distribution list, and re-randomizing the distribution list for re-use once the end of the distribution list is reached during the distribution of the work requests to the distributed processing units as indicated by the entries in the randomized distribution list.

Claim 11-13 (Canceled).

14. (Previously presented) In a data processing network including a network file server and a plurality of virus checking servers, a method comprising:

the network file server obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

the network file server applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

the network file server using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers;

wherein the respective weight for said each virus checking server is programmed into a mapping table, and the network file server indexes the mapping table with said each respective utilization value to obtain the respective weight for said each virus checking server.

Claims 15-16 (Canceled).

17. (Previously presented) In a data processing network including a network file server and a plurality of virus checking servers, a method comprising:

the network file server obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

the network file server applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

the network file server using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers;

wherein the respective weights for the virus checking servers are used for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers by creating a distribution list containing entries indicating the virus checking servers, the respective weight for said each virus checking server specifying the number of the entries indicating said each virus checking server, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the virus checking requests from the network file server to the virus checking servers as indicated by the entries in the randomized distribution list, and re-randomizing the distribution list for re-use once the end of the distribution list is reached during the distributing of the work requests to the virus checking servers as indicated by the entries in the randomized distribution list.

18. (Previously presented) In a data processing network including a network file server and a plurality of virus checking servers, a method comprising:

the network file server obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

the network file server applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

the network file server using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers;

wherein the respective weights for the virus checking servers are used for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers by creating a distribution list containing entries indicating the virus checking servers, the respective weight for said each virus checking server specifying the number of the entries indicating said each virus checking server, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the virus checking requests from the network file server to the virus checking servers as indicated by the entries in the randomized distribution list, and

wherein the network file server obtains the utilization values of the virus checking servers at the start of a heartbeat interval, randomizes the distribution list repetitively during use of the distribution list for load balancing of virus checking requests during the heartbeat interval, obtains new utilization values of the virus checking servers at the start of a following heartbeat

interval, and produces a new distribution list from the new utilization values of the virus checking servers for load balancing of virus checking requests during the following heartbeat interval.

Claims 19-22 (Canceled).

23. (Previously presented) A data processing system comprising distributed processing units and a processor coupled to the distributed processing units for distributing work requests to the distributed processing units, the processor being programmed for:

obtaining a respective utilization value of each distributed processing unit;

applying a function to the respective utilization value of said each distributed processing unit to obtain a respective weight for said each distributed processing unit; and

using the respective weights for the distributed processing units for distributing work requests to the distributed processing units so that the respective weight for said each distributed processing unit specifies a respective frequency at which the work requests are distributed to said each distributed processing unit;

wherein the respective weight for said each distributed processing unit is programmed into a mapping table, and the processor is programmed to apply the function to the respective utilization value of said each distributed processing unit to obtain a respective weight for said each distributed processing unit by indexing the mapping table with said each respective utilization value of said each distributed processing unit to obtain the respective weight for said each distributed processing unit.

Claims 24-27 (Canceled).

28. (Currently amended) A data processing system comprising distributed processing units and a processor coupled to the distributed processing units for distributing work requests to the distributed processing units, the processor being programmed for:

obtaining a respective utilization value of each distributed processing unit;

applying a function to the respective utilization value of said each distributed processing unit to obtain a respective weight for said each distributed processing unit; and

using the respective weights for the distributed processing units for distributing work requests to the distributed processing units so that the respective weight for said each distributed processing unit specifies a respective frequency at which the work requests are distributed to said each distributed processing unit;

wherein the processor is programmed for using the respective weights for the distributed processing units for distributing work requests to the distributed processing units by creating a distribution list containing entries indicating the distributed processing units, the respective weight for said each distributed processing unit specifying the number of the entries indicating said each distributed processing unit, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the work requests to the distributed processing units as indicated by the entries in the randomized distribution list, and re-randomizing the distribution list for re-use once the end of the distribution list is reached during the distribution of the work

requests to the distributed processing units as indicated by the entries in the randomized distribution list.

Claims 29-31 (Canceled).

32. (Previously presented) A data processing system comprising virus checking servers and a network file server coupled to the virus checking servers for distributing virus checking requests to the virus checking servers, the network file server being programmed for:

obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers;

wherein the respective weight for said each virus checking server is programmed into a mapping table, and the network file server is programmed for indexing the mapping table with said each respective utilization value of said each virus checking server to obtain the respective weight for said each virus checking server.

Claims 33-34 (Canceled).

35. (Previously presented) A data processing system comprising virus checking servers and a network file server coupled to the virus checking servers for distributing virus checking requests to the virus checking servers, the network file server being programmed for:

obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers; and

wherein the network file server is programmed for using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers by creating a distribution list containing entries indicating the virus checking servers, the respective weight for said each virus checking server specifying the number of the entries indicating said each virus checking server, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the virus checking requests from the network file server to the virus checking servers as indicated by the entries in the randomized distribution list, and re-randomizing the distribution list for re-use once the end of the distribution list is reached during the distributing of the work requests to the virus checking servers as indicated by the entries in the randomized distribution list.

36. (Previously presented) A data processing system comprising virus checking servers and a network file server coupled to the virus checking servers for distributing virus checking requests to the virus checking servers, the network file server being programmed for:

obtaining a respective utilization value of each virus checking server, the respective utilization value of said each virus checking server indicating a percentage of saturation of said each virus checking server;

applying a function to the respective utilization value of said each virus checking server to obtain a respective weight for said each virus checking server; and

using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers; and

wherein the network file server is programmed for using the respective weights for the virus checking servers for weighted round-robin load balancing of virus checking requests from the network file server to the virus checking servers by creating a distribution list containing entries indicating the virus checking servers, the respective weight for said each virus checking server specifying the number of the entries indicating said each virus checking server, and by randomizing the distribution list, and accessing the randomized distribution list for distributing the virus checking requests from the network file server to the virus checking servers as indicated by the entries in the randomized distribution list, and

wherein the network file server is programmed for collecting utilization statistics from the virus checking servers at the start of a heartbeat interval, for randomizing the distribution list

repetitively during use of the distribution list for load balancing of virus checking requests during the heartbeat interval, for collecting a new set of utilization statistics from the virus checking servers at the start of a following heartbeat interval, and for producing a new distribution list from the new set of utilization statistics for load balancing of virus checking requests during the following heartbeat interval.